D-Wave releases its next-generation quantum annealing chip

reader comments

31 with 28 posters participating, including story author

Share this story

Today, quantum computing company D-Wave is announcing the availability of its next-generation quantum annealer, a specialized processor that uses quantum effects to solve optimization and minimization problems. The hardware itself isn't much of a surprise—D-Wave was discussing its details months ago—but D-Wave talked with Ars about the challenges of building a chip with over a million individual quantum devices. And the company is coupling the hardware's release to the availability of a new software stack that functions a bit like middleware between the quantum hardware and classical computers.

Quantum annealing

Quantum computers being built by companies like Google and IBM are general-purpose, gate-based machines. They can solve any problem and should show a vast acceleration for specific classes of problems—or they will, as soon as the gate count gets high enough. Right now, these quantum computers are limited to a few-dozen gates and have no error correction. Bringing them up to the scale needed presents a series of difficult technical challenges.

D-Wave's machine is not general-purpose; it's technically a quantum annealer, not a quantum computer. It performs calculations that find low-energy states for different configurations of the hardware's quantum devices. As such, it will only work if a computing problem can be translated into an energy-minimization problem in one of the chip's possible configurations. That's not as limiting as it might sound, since many forms of optimization can be translated to an energy minimization problem, including things like complicated scheduling issues and protein structures.

It's easiest to think of these configurations as a landscape with a series of peaks and valleys, with the problem-solving being the equivalent of searching the landscape for the lowest valley. The more quantum devices there are on D-Wave's chip, the more thoroughly it can sample the landscape. So ramping up the qubit count is absolutely critical for a quantum annealer's utility.

This idea matches D-Wave's hardware pretty well, since it's much easier to add qubits to a quantum annealer; the company's current offering has 2,000 of them. There's also a matter of fault tolerance. While errors in a gate-based quantum computer typically result in a useless output, failures on a D-Wave machine usually mean the answer it returns is low-energy, but not the lowest. And for many problems, a reasonably optimized solution can be good enough.

What has been less clear is whether the approach offers clear advantages over algorithms run on classical computers. For gate-based quantum computers, researchers had already worked out the math to show the potential for quantum supremacy. That isn't the case for quantum annealing. Over the last few years, there have been a number of cases where D-Wave's hardware showed a clear advantage over classical computers, only to see a combination of algorithm and hardware improvements on the classical side erase the difference.

Across generations

D-Wave is hoping that the new system, which it is terming Advantage, is able to demonstrate a clear difference in performance. Prior to today, D-Wave offered a 2,000 qubit quantum optimizer. The Advantage system scales that number up to 5,000. Just as critically, those qubits are connected in additional ways. As mentioned above, problems are structured as a specific configuration of connections among the machine's qubits. If a direct connection between any two isn't available, some of the qubits have to be used to make the connection and are thus unavailable for problem solving.

The 2,000 qubit machine had a total of 6,000 possible connections among its qubits, for an average of three for each of them. The new machine ramps up that total to 35,000, an average of seven connections per qubit. Obviously, this enables far more problems to be configured without using any qubits to establish connections. A white paper shared by D-Wave indicates that it works as expected: larger problems fit in to the hardware, and fewer qubits need to be used as bridges to connect other qubits.

Each qubit on the chip is in the form of a loop of superconducting wire called a Josephson junction. But there are a lot more than 5,000 Josephson junctions on the chip. "The lion's share of those are involved in superconducting control circuitry," D-Wave's processor lead, Mark Johnson, told Ars. "They're basically like digital-analog converters with memory that we can use to program a particular problem."

To get the level of control needed, the new chip has over a million Josephson junctions in total. "Let's put that in perspective," Johnson said. "My iPhone has got a processor on it that's got billions of transistors on it. So in that sense, it's not a lot. But if you're familiar with superconducting integrated circuit technology, this is way on the outside of

the curve.” Connecting everything also required over 100 meters of superconducting wire—all on a chip that’s roughly the size of a thumbnail.

While all of this is made using standard fabrication tools on silicon, that’s just a convenient substrate—there are no semiconducting devices on the chip. Johnson wasn’t able to go into details on the fabrication process, but he was willing to talk about how these chips are made more generally.

This isn’t TSMC

One of the big differences between this process and standard chipmaking is volume. Most of D-Wave’s chips are housed in its own facility and get accessed by customers over a cloud service; only a handful are purchased and installed elsewhere. That means the company doesn’t need to make very many chips.

When asked how many it makes, Johnson laughed and said, “I’m going to end up as the case of this fellow who predicted there would never be more than five computers in this world,” before going on to say, “I think we can satisfy our business goals with on the order of a dozen of these or less.”

If the company was making standard semiconductor devices, that would mean doing one wafer and calling it a day. But D-Wave considers it progress to have reached the point where it’s getting one useful device off every wafer. “We’re constantly pushing way past the comfort zone of what you might have at a TSMC or an Intel, where you’re looking for how many 9s can I get in my yield,” Johnson told Ars. “If we have that high of a yield, we probably haven’t pushed hard enough.”

A lot of that pushing came in the years leading up to this new processor. Johnson told Ars that the higher levels of connectivity required a new process technology. “[It’s] the first time we’ve made a significant change in the technology node in about 10 years,” he told Ars. “Our fab cross-section

is much more complicated. It’s got more materials, it’s got more layers, it’s got more types of devices and more steps in it.”

Beyond the complexity of fashioning the device itself, the fact that it operates at temperatures in the milli-Kelvin range adds to the design challenges as well. As Johnson noted, every wire that comes in to the chip from the outside world is a potential conduit for heat that has to be minimized—again, a problem that most chipmakers don’t face.

https://arstechnica.com/?p=1710391